Submitted to the Annals of Applied Statistics

MODELING AND ESTIMATION FOR SELF-EXCITING SPATIO-TEMPORAL MODELS OF TERRORIST ACTIVITY

By Nicholas J. Clark* and Philip M. Dixon

Iowa State University

arXiv:1703.08429v1 [stat.AP] 24 Mar 2017

Spatio-temporal hierarchical modeling is an extremely attractive way to model the spread of crime or terrorism data over a given region, especially when the observations are counts and must be modeled discretely. The spatio-temporal diffusion is placed, as a matter of convenience, in the process model allowing for straightforward estimation of the diffusion parameters through Bayesian techniques. However, this method of modeling does not allow for the existence of self-excitation, or a temporal data model dependency, that has been shown to exist in criminal and terrorism data. In this manuscript we will use existing theories on how violence spreads to create models that allow for both spatio-temporal diffusion in the process model as well as temporal diffusion, or self-excitation, in the data model. We will further demonstrate how Laplace approximations similar to their use in Integrated Nested Laplace Approximation can be used to quickly and accurately conduct inference of self-exciting spatio-temporal models allowing practitioners a new way of fitting and comparing multiple process models. We will illustrate this approach by fitting a self-exciting spatio-temporal model to terrorism data in Iraq and demonstrate how choice of process model leads to differing conclusions on the existence of self-excitation in the data and differing conclusions on how violence is spreading spatio-temporally.

1. Introduction. A typical spatio-temporal model consists of three levels, a data model, a process model, and a parameter model. A common way to model data then is to assume $Y(.)$, is conditionally independent given the process model $X(.)$. For example, if observations take place at aerial regions, $\boldsymbol{s_i}$, at discrete time periods, t, and $Y(\boldsymbol{s_i}, t)$ are counts, a common model is $Y(\boldsymbol{s_i}, t)|X(\boldsymbol{s_i}, t) \sim$ indep $Pois(\mu(\boldsymbol{s_i}, t))$. The spatio-temporal diffusion structure is commonly then placed on the process model which commonly is assumed to have a Gaussian joint distribution of $\boldsymbol{X} \sim$ Gaus $(\boldsymbol{0}, \Sigma(\theta))$. The majority of analysis of these models is done using Bayesian techniques requiring a further parameter model for θ. The challenge in these models is, then, the structure of $\Sigma(\theta)$. If both the covariance and the precision is

*Supported in Part through a Omar Nelson Bradley Fellowship

MSC 2010 subject classifications: Primary 62H11; secondary 62P25

Keywords and phrases: Laplace Approximations, non-seperable space-time, spatial-temporal dependence, Bayesian, terrorism

chosen to be too dense inference quickly becomes impossible as the size of $\Sigma(\theta)$ is typically, for spatio-temporal models, larger than 10^4 x 10^4. A thorough overview giving many examples of this method of modeling is given in Cressie and Wikle (2012).

In the statistical modeling of terrorism or crime data one modeling choice, is to use an extremely general spatio-temporal process model to capture variance not explained through the use of covariates. For example Python et al. (2016) use a Matern class covariance function over space and an AR(1) process and then use covariates to test the impact of infrastructure, population, and governance. The general spatio-temporal process models used, in this case, has an extremely sparse precision structure greatly aiding in computations.

While diffusion in spatio-temporal models is often modeled through a latent process, more recent models describing the spread of violence have incorporated self-excitation, or spatio-temporal diffusion that exists linearly in the data model itself. Self-excitation is the theory that in terrorism, or crime, the probability of an event occurring is a function of previous successful events. For instance Mohler et al. (2012) demonstrate that burglars are more likely to rob locations that have previously, successfully, been robbed. Mohler et al. (2013) derived a class of models that allowed for temporal diffusion in both the process model as well as the data model and demonstrated how the two processes were identifiable.

In the modeling of terrorism data Lewis et al. (2012), Porter et al. (2012), and Mohler et al. (2013) have all successfully used the self-excitation approach to model. Most recently, Tench, Fry and Gill (2016) used a likelihood approach for temporal modeling of IEDs in Northern Ireland using self-excitation. However, in these papers, the existence and analysis of self-excitation was the primary objective and any process model dependency was either ignored or treated as a nuisance. The one exception is in Mohler et al. (2013) where a temporal only model was assumed for the process model and inference was conducted on both the process model dependency and the data model dependency.

In this manuscript, we will consider a spatial and a spatial-temporal process model that allows for self-excitation. We will present two self-exciting models for terrorist activity that have different process models corresponding to different notions of how terrorism evolves in time and space as well as temporal dependency in the data model to account for self-excitation. These two models are specific cases of more general spatio-temporal models that allow dependency in both the process model as well as the data model.

We will further show how Laplace approximations similar to their use

in Integrated Nested Laplace Approximation, or INLA, an approximate Bayesian method due to Rue, Martino and Chopin (2009) can be used to conduct inference for these types of models. We will show, via simulation, how INLA, when appropriately modified, can accurately be used to make inference on process level parameters for self-exciting models and aid analysts in determining the appropriate process model when scientific knowledge cannot be directly applied as in Cressie and Wikle (2012). Finally, we will apply this technique to terrorism data in Iraq. We will show that choice of process model, in this case, results in differing conclusions on the impact of self-excitation in the model.

2. Self-Exciting Spatio-Temporal Models. The use of self-exciting models in both criminal and terrorism modeling has become increasingly popular over the last decade after being originally introduced in Short et al. (2008). Self-excitement, in a statistical model, directly models copy-cat behavior by letting an observed event increase the intensity (or excites a model) over a specified time or location. Self-exciting models are closely related to Hawkes processes, which are counting processes where the probability of an event occurring is directly related to the number of events that previously occurred. In a self-exciting model, the criminal intensity at a given spatio-temporal location, (x, y, t) is a mixture of a background rate, ν and self excitement function, $f(\mathcal{H}_{x,y,t})$ that is dependent on the observed history at that location, $\mathcal{H}_{x,y,t}$.

A common temporal version of a Hawkes process is

$$\begin{aligned} Y(t_i) &\sim Pois(\lambda_i) \\ \lambda_i &= \nu + \sum_{j<i} \kappa(i-j) y(t_j) \end{aligned} \tag{2.1}$$

In this example, in order for the process to have finite expectation in the limit, $\kappa(i-j)$ must be positive and $\sum_{i=1}^{\infty} \kappa(i) < 1$. $\kappa(i-j)$ can be thought of as the probability that an event at time j triggers an event at time i.

Laub, Taimre and Pollett (2015) provides an excellent overview of the mathematical properties of the continuous Hawkes process and the discrete process when $\kappa(i-j)$ is taken to be an exponential decay function.

In Mohler et al. (2012), ν was treated as separable in space and time and was non-parametrically estimated using stochastic declustering, while in Mohler et al. (2013), the spatial correlations were ignored and an AR(1) process was used for $\kappa(t)$ and an exponential decay was assumed for the self-excitation. In terrorism modeling Lewis et al. (2012) used a piece-wise linear function for $\kappa(t)$.

Here we will first define a general model that allows spacial or spacial-temporal correlation to exist in the process model and positive temporal correlation to exist in the data model to allow for self-excitation. First define $s_i \in \mathcal{R}^2$, $i \in \{1, 2, ..., s\}$ as locations in a fixed, aerial, region. We further define $t \in \{1, 2, ..., n\}$ as discrete time. The general form of a spacial-temporal self-exciting model is then given in (2.2)

$$
\begin{aligned}
&Y(\boldsymbol{s_i}, t)|\mu(\boldsymbol{s_i}, t) \sim Pois(\mu(\boldsymbol{s_i}, t)) \\
&\mu(\boldsymbol{s_i}, t) = \exp(X(\boldsymbol{s_i}, t)) + \eta Y(\boldsymbol{s_i}, t-1) \\
&\boldsymbol{X} \sim Gau(\mathbf{0}, \Sigma(\boldsymbol{\theta}))
\end{aligned} \tag{2.2}
$$

Comparing the above to the Hawkes process, we now have ν as a function of space-time and denote it as $X(\boldsymbol{s_i}, t)$. We use the simplest form of self-excitation letting $\kappa(i-j)$ be a point-mass function such that $\kappa(k) = \eta$ for $k = 1$ and $\kappa(k) = 0$ for $k \neq 1$. In all cases $Y(\boldsymbol{s_i}, t)$ will be discrete, observable, count data.

To contrast (2.2) with a typical spatial model, figure 1 depicts the expectation for one aerial location $(\boldsymbol{s_i}, t)$ without self-excitation and with self-excitation as shown in Figure 1. In this figure, the lower line shows $\mu(\boldsymbol{s_i}, t)$ with $\eta = 0$, and the upper line has $\eta = .4$. The impact of self-excitation is clearly present in time 10-13.

2.1. *Spatially Correlated Self-Exciting Model.* In the first example of (2.2) we assume the background intensity rate, $X(\boldsymbol{s_i}, t)$ has only spatial correlation. This model is motivated through the assumption that the latent dependency, $X(\boldsymbol{s_i}, t)$, is as a continuous measure of violent tendency at region s_i at time period t and regions that are closer together in space are assumed to share common characteristics.

Next, define $N(\boldsymbol{s_i})$ as the neighborhood of location s_i where two regions are assumed to be neighbors if they share a common border. $|N(\boldsymbol{s_i})|$ is the number of neighbors of location s_i. The model for $Y(\boldsymbol{s_i}, t)$, or the number of observed violent events at a given space-time location is then given by:

$$
\begin{aligned}
&Y(\boldsymbol{s_i}, t)|\mu(\boldsymbol{s_i}, t) \sim Pois(\mu(\boldsymbol{s_i}, t)) \\
&\mu(\boldsymbol{s_i}, t) = \exp(X(\boldsymbol{s_i}, t)) + \eta Y(\boldsymbol{s_i}, t-1) \\
&X(\boldsymbol{s_i}, t) = \theta_1 \sum_{\boldsymbol{s_j} \in N(\boldsymbol{s_i})} X(\boldsymbol{s_j}, t) + \epsilon(\boldsymbol{s_i}, t) \\
&\epsilon(\boldsymbol{s_i}, t) \sim Gau(0, \sigma^2)
\end{aligned} \tag{2.3}
$$

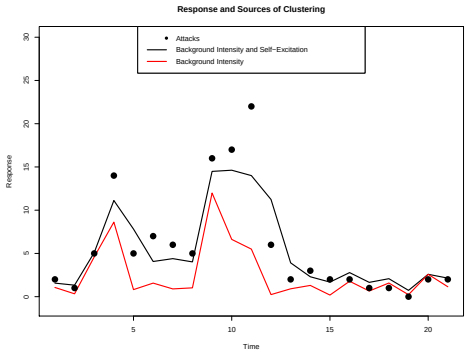

FIG 1. *This figure shows an example of the expectation of two processes, one with self-excitation and one without. The bottom line is the expectation of a process withs no self-excitation, the top has self-excitation of* $\eta = .4$. *The data realizations are from the process with self-excitation.*

Letting $\boldsymbol{H}$ denote the spatial neighborhood matrix such that $H_{i,j} = H_{j,i} = 1$ if $\boldsymbol{s_i}$ and $\boldsymbol{s_j}$ are neighbors, the full distribution of the joint distribution of the latent state is $\boldsymbol{X} \sim Gau(\boldsymbol{0}, (\boldsymbol{I}_{ns,ns} - \boldsymbol{I}_{n,n} \otimes \theta_1 \boldsymbol{H})^{-1} \boldsymbol{L} (\boldsymbol{I}_{ns,ns} - \boldsymbol{I}_{n,n} \otimes \theta_1 \boldsymbol{H})^{-1})$ where $\boldsymbol{L} = \text{diag}(\sigma^2, ..., \sigma^2)$. The evolution in the latent field is equivalent to the spatial evolution in what is commonly referred to as a Simultaneous or Spatial Auto-regressive model (SAR). Alternatively, the Conditional Auto-regressive model (CAR) of Besag (1974) could be used to model the latent state modifying the joint density above.

The difference between the SAR and (2.3) is in the self excitement parameter, η. In (2.3), temporal correlation is present, but is present through the data model specification rather than through a temporal evolution in the latent state. Therefore, temporal correlation is a function solely of the self-excitation in the process. η gives the probability of an event at time $t-1$ creating an event at time t. In order for the system to be well-behaved, η is constrained to (0,1). In order for the joint distribution of $\boldsymbol{X}$ to be valid, $\theta_1 \in (\psi_{(1)}^{-1}, \psi_{(n)}^{-1})$ where ψ_i is the ith smallest eigenvalues of $\boldsymbol{H}$.

The critical assumption in this model is that the propensity of a given location to be violent is spatially correlated with its adjacent spatial neigh-

bors and only evolves over time as a function of excitation. If terrorism is diffusing according to this model, regions that are geographically adjacent are behaving in a similar manner. The existence of self-excitation would indicate that individuals within a region are being inspired through the actions of others. While combating terrorism is complex, if terrorism is diffusing in this manner, one suggestion would be to identify the root causes within a geographic area as well as quick action against any malicious actor to discourage copy-cat behavior.

2.2. *Reaction Diffusion Self-Exciting Model.* Alternatively, a model similar to Short et al. (2008) can be used to motivate the process model for the latent state resulting in a non-separable spatio-temporal, $\boldsymbol{X}$. Here we let $X(\boldsymbol{s_i}, \boldsymbol{t})$ corresponds to a continuous measure of violence due to terrorists or criminals at location $\boldsymbol{s_i}$ at time t. This is still a latent variable as we are not directly measuring $X(\boldsymbol{s_i}, \boldsymbol{t})$. However, now in order for an area to increase in violent tendency, a neighboring area must decrease as the actors causing the violence move from location to location. Furthermore, if terrorists are removed from the battlefield at a rate proportional to the total number of terrorists present and if terrorists move to fill power vacuums, the process model is similar to the reaction-diffusion partial differential equation (see Cressie and Wikle (2012) for more on the reaction-diffusion model)

$$\frac{\partial X(\boldsymbol{s_i}, t)}{\partial t} = \frac{\beta}{|N(\boldsymbol{s_i})|} \triangle X(\boldsymbol{s_i}, t) - \alpha X(\boldsymbol{s_i}, t) \tag{2.4}$$

In order to generalize this PDE to an irregular lattice, we make use of the graphical Laplacian, Γ, in place of $\triangle$ in (2.4). Γ is a matrix that extends the notion of second derivatives to irregular graphs and can be defined as a matrix of the same dimension as the number of geographical regions with entries given by

$$\Gamma(s_i, s_j) \begin{cases} -N(s_i) & j = i \\ 1 & j \in N(s_i) \\ 0 & \text{Otherwise} \end{cases}$$

With the addition of a random noise term assumed to be Gaussian, the full model can be seen as an example of (2.2).

$$
\begin{aligned}
(2.5)\quad & Y(\boldsymbol{s_i},t)|\mu(\boldsymbol{s_i},t) \sim Pois(\mu(\boldsymbol{s_i},t)) \\
& \mu(\boldsymbol{s_i},t) = \exp(X(\boldsymbol{s_i},t)) + \eta Y(\boldsymbol{s_i},t-1) \\
& X(\boldsymbol{s_i},t) = \frac{\beta}{|N(s_i)|} \sum_{\boldsymbol{s_j}\in N(\boldsymbol{s_i})} X(\boldsymbol{s_j},t-1) + (1-\beta-\alpha)X(\boldsymbol{s_i},t-1) + \epsilon(\boldsymbol{s},t) \\
& \epsilon(\boldsymbol{s},t) \sim Gau(0,\sigma^2)
\end{aligned}
$$

In contrast to the Spatially Correlated Self-Exciting (SCSE) Model, the process model dependency exists in both space and time. In order to derive properties of this model we first let $\boldsymbol{M} = -\beta \text{ diag}\left(\frac{1}{|N_{s_i}|}\right)\Gamma + (1-\alpha)\boldsymbol{I}_{s,s}$ and now note that this is equivalent to a Vector Auto-Regressive, VAR, model $\boldsymbol{X}_t = \boldsymbol{M}\boldsymbol{X}_{t-1} + \boldsymbol{\epsilon}$ with $\boldsymbol{\epsilon} \sim Gau(\boldsymbol{0}, \sigma^2\boldsymbol{I})$.

The VAR(1) model requires all the eigenvalues of $\boldsymbol{M}$ to be between -1 and 1. This can be satisfied by first noting that 0 is always an eigenvalue of $\text{diag}\left(\frac{1}{|N_{s_i}|}\right)$ trivially corresponding to the eigenvector of all 1s. The largest eigenvalue is at most 2 as shown in Chung (1997). Due to the structure of $(1-\alpha)\boldsymbol{I}_{s,s}$ this implies maximum eigenvalue of $\boldsymbol{M}$ is $(1-\alpha)$ and minimum is $-2\beta + (1-\alpha)$. Therefore, the parameter spaces for α and β are $\alpha \in (0,1)$ and $\beta \in (\frac{-\alpha}{2}, \frac{2-\alpha}{2})$.

Just as in the SCSE Model, if ϵ has a Gaussian distribution, the Reaction Diffusion Self-Exciting (RDSE) Model has an exact solution for the latent Gaussian field, $\boldsymbol{X}$.

Letting Σ_s be the spatial covariance at a fixed period of time which is assumed to be invariant to time , then we can solve for Σ_s by using the relationship $\Sigma_s = \boldsymbol{M}\Sigma_s\boldsymbol{M^T} + \sigma^2\boldsymbol{I}$. As demonstrated by Cressie and Wikle (2012), this leads to $\text{vec}(\Sigma_s) = \left(\boldsymbol{I}_{s^2,s^2} - \boldsymbol{M}\otimes\boldsymbol{M}\right)^{-1}\text{vec}\left(\sigma^2\boldsymbol{I}_{s,s}\right)$ where vec () is the matrix operator that stacks each column of the matrix on top of one or another. Recall that s is the size of the lattice that is observed at each time period. The joint distribution for all $\boldsymbol{X}$ is then $\boldsymbol{X} \sim Gau(\boldsymbol{0}, \boldsymbol{\Sigma})$ where

$$
(2.6)\qquad \Sigma = \left[\begin{array}{c|c|c|c} \Sigma_s & M\Sigma_t & \ldots & M^n\Sigma_s \\ \hline M\Sigma_s & \Sigma_s & \ldots & M^{n-1}\Sigma_s \\ \hline \ldots & \ldots & \ldots & \ldots \\ \hline M^n\Sigma_s & M^{n-1}\Sigma_s & \ldots & \Sigma_s \end{array}\right]
$$

However, practically, this involves inverting a potentially large matrix $\boldsymbol{I}_{s^2,s^2} - \boldsymbol{M}\otimes\boldsymbol{M}$. Therefore, it is easier to deal with the inverse of (2.6)

given in (2.7).

(2.7)

$$\Sigma^{-1} = \left[\begin{array}{c|c|c|c|c} \boldsymbol{I}_{n,n} & -M & \mathbf{0} & \ldots & \ldots \\ \hline -M^T & M^TM + \boldsymbol{I}_{n,n} & -M & \mathbf{0} & \ldots \\ \hline \mathbf{0} & -M^T & M^TM + \boldsymbol{I}_{n,n} & -M & \ldots \\ \hline \ldots & \ldots & \ldots & \ldots & \ldots \\ \hline \mathbf{0} & \ldots & -M^T & M^TM + \boldsymbol{I}_{n,n} & -M \\ \hline \mathbf{0} & \ldots & \ldots & -M^T & \boldsymbol{I}_{n,n} \end{array}\right] \frac{1}{\sigma^2}$$

The primary difference between the SCSE model and the RDSE model is that the process model correlation in the SCSE is only spatial while in the RDSE it is spatio-temporal. In the below toy examples, we show the expectation for $X(s_i, t)$ for both the SCSE and the RDSE model on a 4 x 4 lattice structure. We fixed both models with a value of $X(s_1, 1) = 10$ as the upper left hand observation at time 1. As seen in the RDSE model, the high count at time 1 spreads to neighboring regions in time 2 and time 3 whereas the process model has no temporal spread in the SCSE but has a high level of spatial spread.

Spatially Correlated Latent Process Conditional on $(s_1, 1) = 10$

Time 1				Time 2			
10	5	2	1	0	0	0	0
5	4	3	2	0	0	0	0
2	3	2	1	0	0	0	0
1	1	1	1	0	0	0	0

Reaction Diffusion Latent Process Conditional on $(s_1, 1) = 10$

Time 1				Time 2				Time 3			
10	1	1	0	3	3	0	0	2	1	1	0
1	1	0	0	3	1	0	0	1	1	0	0
1	0	0	0	0	0	0	0	1	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0

Practically, if data follows the RDSE model, it implies a high terrorism count in one region will manifest into a high terrorism count in a neighboring region at a later time period. In combating terrorism, the RDSE might suggest isolating geographical regions to mitigate the risk of spread while addressing self-excitation through direct action against malicious actors who are inspiring others.

3. Model Fitting. In both the RDSE and the SCSE, spatio-temporal diffusion exists in both the process model and the data model. If the diffusion was solely in the process model, a technique for inference would be Integrated Nested Laplace Approximation, or INLA.

INLA was first proposed in Rue, Martino and Chopin (2009) to specifically address the issue of Bayesian Inference of high dimensional Latent Gaussian Random Fields, LGRFs. An example of this for count data is:

$$\begin{aligned} Y(s_i) &\sim Pois(\mu(s_i,t)) \\ \mu(\boldsymbol{s_i},\boldsymbol{t}) &= \exp(\lambda(\boldsymbol{s_i},t)) \\ \lambda(\boldsymbol{s_i},t) &= \beta_0 + \boldsymbol{Z}^t\boldsymbol{\beta} + X(\boldsymbol{s_i},\boldsymbol{t}) \\ X(\boldsymbol{s_i},\boldsymbol{t}) &\sim Gau(\mathbf{0},\Sigma(\theta)) \end{aligned} \tag{3.1}$$

INLA is often preferable over MCMC for these types of models. An issue with traditional Markov Chain Monte Carlo (MCMC) techniques for these models is that the dimension of X is often very large. Therefore, while MCMC has $O_p(N^{-1/2})$ errors, the N in the errors is the simulated sample size for the posterior. Just getting $N = 1$ may be extremely difficult due to the vast number of elements of X that need to be estimated. In general, MCMC will take hours or days in order to successfully simulate from the posterior making the computational cost of fitting multiple process models extremely high. In Python et al. (2016), terrorism data was fit using a grid over the entire planet using INLA, though without self-excitation in the model.

To address the issues with MCMC use in LGRFs, Rue, Martino and Chopin (2009) developed a deterministic approach based on multiple Laplacian approximations. Using the general notation for LGRFs

$$\pi(\boldsymbol{\theta},\boldsymbol{X}|\boldsymbol{Y}) \propto \pi(\theta)|Q(\boldsymbol{\theta})|^{1/2}\exp\left[\frac{-1}{2}\boldsymbol{X}^tQ(\boldsymbol{\theta})\boldsymbol{X} + \sum_{\boldsymbol{s}}\log\left(\pi(Y(\boldsymbol{s_i})|X(\boldsymbol{s_i}),\boldsymbol{\theta})\right)\right] \tag{3.2}$$

In order to conduct inference on this model, we need to estimate $\pi(\boldsymbol{\theta}|\boldsymbol{y})$, $\pi(\theta_i|\boldsymbol{y})$ and $\pi(x_i|\boldsymbol{y})$. The main tool Rue, Martino and Chopin (2009) employ is given in their equation (3) as

$$\tilde{\pi}(\boldsymbol{\theta}|\boldsymbol{Y}) \propto \frac{\pi(\boldsymbol{X},\boldsymbol{\theta},\boldsymbol{Y})}{\tilde{\pi}_G(\boldsymbol{X}|\boldsymbol{\theta},\boldsymbol{Y})}|_{X=x^*(\boldsymbol{\theta})} \tag{3.3}$$

In Rue, Martino and Chopin (2009) they note that the denominator of (3.3) almost always appears to be unimodal and approximately Gaussian. The authors then propose to use a Gaussian approximation to $\pi(\boldsymbol{X}|\boldsymbol{\theta},\boldsymbol{Y})$

which is denoted above as $\tilde{\pi}_G$. Moreover, (3.3) should hold no matter what choice of $\boldsymbol{X}$ is used, so a convenient choice for $\boldsymbol{X}$ is the mode for a given θ, which Rue, Martino and Chopin (2009) denote as $x^*(\boldsymbol{\theta})$.

Now, $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ can be explored by calculating the marginal for choices of θ, which if chosen carefully can greatly decrease the computational time. These explored values can then be numerically integrated out to get credible intervals for $\pi(\theta_i|\boldsymbol{Y})$.

Following the exploration of $\theta|Y$, and computation of $\theta_i|Y$, INLA next proceeds to approximate $\pi(X(s_i)|\boldsymbol{\theta}, \boldsymbol{Y})$. The easiest way to accomplish this is to use the marginals that can be derived straightforwardly from $\tilde{\pi}_G(\boldsymbol{X}|\boldsymbol{\theta}, \boldsymbol{Y})$ from (3.3). In this manuscript we will use this technique for simplicity of computation, however, if the latent states are of interest in the problem (and they often are), this can be problematic as it fails to capture any skewness of the posterior of $\boldsymbol{X}$. One way to correct this is to re-apply (3.3) in the following manner:

$$\tilde{\pi}_{LA}(X(\boldsymbol{s_i})|\theta, y) \propto \frac{\pi(\boldsymbol{X}, \theta, \boldsymbol{Y})}{\tilde{\pi}_G(\boldsymbol{X}_{-s_i}|X(s_i), \boldsymbol{\theta}, \boldsymbol{Y})}|_{x_{-i}=\boldsymbol{x_{-i}}^*(x_i,\theta)} \tag{3.4}$$

In (3.4) $\boldsymbol{X}_{-s_i}$ is used to represent $\boldsymbol{X}$ with latent variable $X(s_i)$ removed. This is a reapplication of Tierney and Kadane's marginal posterior density and gives rise to the nested term in INLA.

3.1. *Laplace Approximation for Spatio-Temporal Self-Exciting Models.* While INLA is an attractive technique due to computational speed and implementation, it is not immediately usable for the SCSE and the RDSE as the structure in (2.2) is

$$\begin{aligned} \mu(\boldsymbol{s_i}, t) &= \exp(X(\boldsymbol{s_i}, t)) + \eta Y(\boldsymbol{s_i}, t-1) \\ \eta &\in (0, 1) \end{aligned} \tag{3.5}$$

In this structure, $X(.)$ and $Y(.)$ are not linearly related and a Gaussian prior for η is clearly not appropriate due to the parameter space constraints.

However, Laplace approximations can still be used by conducting inference on η at the same time inference is conducted on the the set of latent model parameters. In both the Spatially Correlated Self-Exciting Model and the Reaction Diffusion Self-Exciting model, the full conditional for the latent state is

$$\pi(\boldsymbol{X}|\boldsymbol{Y}, \boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\boldsymbol{X}^T\boldsymbol{Q}(\boldsymbol{\theta})\boldsymbol{X} + \sum_{s_i,t} \log \pi\left(Y(\boldsymbol{s_i}, t)|X(\boldsymbol{s_i}, t), \eta, Y(\boldsymbol{s_i}, t-1)\right)\right) \tag{3.6}$$

Here we will let $\boldsymbol{\theta} = (\theta_1, \sigma^2, \eta)^T$ and $\boldsymbol{Q}(\boldsymbol{\theta}) = (\boldsymbol{I}_{nt,nt} - \theta_1 \boldsymbol{I}_{t,t} \otimes \boldsymbol{H})$ for the Spatially Correlated Self-Exciting Model and $\boldsymbol{\theta} = (\alpha, \beta, \sigma^2, \eta)^T$ and $\boldsymbol{Q} =$

$\boldsymbol{\Sigma^{-1}}$ in the Reaction-Diffusion Self-Exciting Model where Σ is as in (2.6). While $\boldsymbol{\theta}$ in (3.6) does not contain η we next do a Taylor series expansion of $\log \pi\left(Y(\boldsymbol{s_i}, t) | X(\boldsymbol{s_i}, t), \eta, Y(\boldsymbol{s_i}, t-1)\right)$, as a function of $X(\boldsymbol{s_i}, t)$ and, for each $\boldsymbol{s_i}, t$, expand the term about a guess for the mode, say $\mu_0(\boldsymbol{s_i}, \boldsymbol{t})$. Now, defining $\boldsymbol{\mu_0}$ as the vector of all $\mu_0(\boldsymbol{s_i}, \boldsymbol{t})$ we have:

$$\pi(\boldsymbol{X}|\boldsymbol{Y}, \boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\boldsymbol{X}^T\left(\boldsymbol{Q^*(\theta)}|\boldsymbol{\mu_0}\right)\boldsymbol{X} + \boldsymbol{X}^T\left(\boldsymbol{B^*}(\boldsymbol{\theta})|\boldsymbol{\mu_0}\right)\right) \tag{3.7}$$

Where $\boldsymbol{Q^*(\theta)}|\mu_0$ is an updated precision matrix.

$$\boldsymbol{Q^*(\theta)}|\boldsymbol{\mu_0} = \boldsymbol{Q(\theta)} + \text{diag}\left(-\frac{\partial^2 \log \pi\left(Y(\boldsymbol{s_i}, t)\right)}{\partial X(\boldsymbol{s_i}, t)^2}\right)\Big|_{X(\boldsymbol{s_i},t)=\mu(\boldsymbol{s_i},t)} \tag{3.8}$$

While in (3.8), $\boldsymbol{Q}(\boldsymbol{\theta})$ does not contain η, $\boldsymbol{Q^*(\theta)}$ does depend on the self-excitation parameter.

We can further define $\boldsymbol{B^*}(\boldsymbol{\theta}|\mu_0)$ as a vector of the same length as $X(\boldsymbol{s_i}, t)$ where each element is given by

(3.9)

$$B(\boldsymbol{s_i}, t|\mu_0) = \left(\frac{\partial \log \pi\left(Y(\boldsymbol{s_i}, t)\right)}{\partial X(\boldsymbol{s_i}, t)}\Big|_{X(\boldsymbol{s_i},t)=\mu(\boldsymbol{s_i},t)} - \mu(\boldsymbol{s_i}, t)\frac{\partial^2 \log \pi\left(Y(\boldsymbol{s_i}, t)\right)}{\partial X(\boldsymbol{s_i}, t)^2}\Big|_{X(\boldsymbol{s_i},t)=\mu(\boldsymbol{s_i},t)}\right)$$

Now, we find the values of $\mu(\boldsymbol{s_i})$ that maximize (3.7). This is done through the use of an iterative maximization algorithm by solving for $\boldsymbol{\mu_1}$ in $\left(Q^*(\boldsymbol{\theta})|\boldsymbol{\mu_0}\right)\boldsymbol{\mu_1} = \boldsymbol{B^*}(\boldsymbol{\theta}|\mu_0)$. For a fixed $\boldsymbol{\theta}$, this converges rapidly, usually within 5-6 iterations.

In (3.3), for a fixed $\boldsymbol{\theta}$, we can then find $x^*(\boldsymbol{\theta})$. When the denominator of (3.3) is evaluated at $x^*(\boldsymbol{\theta})$ it becomes $|\boldsymbol{Q^*(\theta)}\frac{1}{2\pi}|^{1/2}$ which is equivalent to the hyperparameter inference recommended by Lee and Nelder (1996) as pointed out by R. A. Rigby in Rue, Martino and Chopin (2009).

In order to best explore $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ the posterior mode is first found through either a gradient descent or Newton-Raphson based method. While the Newton-Raphson converges quicker, it requires multiple computations of a numerical Hessian, which oftentimes requires a steeper computational price than the simpler gradient methods. In our simulations described below we found a gradient based method with step-halving worked relatively quickly.

After locating the posterior mode of $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$, the parameter space can be explored using the exploration strategy laid out in section 3.1 of Rue, Martino and Chopin (2009).

Now, for the set of diffusion parameters, $\boldsymbol{\theta}$ which contain η, we have a method of estimating $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$. Inference for any further data model covariates can now be conducted in the same manner as done in Rue, Martino and Chopin (2009).

3.2. *Model Comparison and Goodness of Fit.* In order to conduct model comparison, we will use the deviance information criterion (DIC) originally proposed by Spiegelhalter et al. (2002). Goodness of fit will be conducted through the use of posterior predictive p-values, outlined by Gelman, Meng and Stern (1996).

To approximate the DIC, we first find the effective number of parameters for a given $\boldsymbol{\theta}$. As noted in Rue, Martino and Chopin (2009), we can estimate this by using $n - \text{tr}\left(\boldsymbol{Q}(\boldsymbol{\theta})\boldsymbol{Q}^*(\boldsymbol{\theta})^{-1}\right)$ where $\boldsymbol{Q}^*(\boldsymbol{\theta})$ is given in (3.8) and $\boldsymbol{Q}(\boldsymbol{\theta})$ is the process model precision matrix evaluated at $\boldsymbol{\theta}$. This gives the effective number parameters for a given $\boldsymbol{\theta}$, which can then be averaged over $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ to get the effective number of parameters for the model.

Secondly, we calculate the deviance of the mean

$$-2\sum_{s_i,t}\log\pi\left(Y(s_i,t)|\hat{X}(s_i,t),\boldsymbol{\theta^*}\right) \tag{3.10}$$

where $\boldsymbol{\theta^*}$ is the posterior mode and $\hat{X}(s_i,t)$ is the expectation of the latent state fixing $\theta = \theta^*$. DIC can then be found through deviance of the mean plus two times the effective number of parameters as in chapter 7 of Gelman et al. (2014) and initially recommended by Spiegelhalter et al. (2002).

In order to assess goodness of fit in analyzing the terrorism data in Section 5, we will use posterior predictive P-values as described by Gelman, Meng and Stern (1996). Here, we pick critical components of the original dataset that we wish to see if the fitted model can accurately replicate, for instance the number of zeros in the dataset which we can designate as $T(\boldsymbol{Y})$. Next, for an index $m = 1...M$, We then draw a value of $\boldsymbol{\theta_m}$ according to $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ and simulate a set of observations $Y^*(\boldsymbol{s_i}, \boldsymbol{t})_m$ of the same dimension as $\boldsymbol{Y}$ and compute $T(\boldsymbol{Y}^*_m)$. This process is repeated M times and a posterior predictive p-value is computed as $\frac{1}{M}\sum_{m=1}^{M} I\left[T(\boldsymbol{Y}^*_m) > T(\boldsymbol{Y})\right]$ where $I\left[.\right]$ is the indicator function. While not a true P-value, both high and low values of the posterior predictive p-value should cause concern over the fitted models ability to replicate features of the original dataset.

4. Simulation. In order to validate the Laplace based methodology for spatially correlated self-exciting models we conducted simulation studies using data on a 8 by 8 Spatial grid assuming a rook neighborhood structure. In order to decrease the edge effect, we wrapped the grid on a torus so each node had four neighbors. For each grid location we simulated 100 observations, creating a spatio-temporal model that had 6400 observations, meaning in (3.3), $\Sigma(\theta)$ had a dimension of 6400 x 6400.

In the first simulation we used (2.3) fixing the parameters at values that generated data that appeared to resemble the data from Iraq used in Section 5. The generating model we used was:

$$
\begin{aligned}
(4.1) \qquad & Y(\boldsymbol{s_i},t)|\mu(\boldsymbol{s_i},t) \sim \text{Pois }(\mu(\boldsymbol{s_i},t)) \\
& \mu(\boldsymbol{s_i},t) = \exp(-1 + X(\boldsymbol{s_i},t)) + .2Y(\boldsymbol{s_i},t-1) \\
& X(\boldsymbol{s_i},t) = .22 \sum_{\boldsymbol{s_j} \in N(\boldsymbol{s_i})} X(\boldsymbol{s_j},t) + \epsilon(\boldsymbol{s_i},t) \\
& \epsilon(\boldsymbol{s_i},t) \sim Gau(0,.4)
\end{aligned}
$$

The spatial parameter for model was $\theta_1 = .22$ which suggests a positive correlation between spatially adjacent locations. An η value of .2 would suggest that an observed event has a probability of .2 of 'exciting' an event at the next time period. σ^2 was allowed to be .4. Finally, an intercept value of $\beta_0 = -1$ was used to correspond to the fact that in real world applications the latent process likely is not zero mean.

Once the data were generated, we found $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ by applying (3.3). Here we note that the numerator of (3.3) is $\pi(\boldsymbol{X},\boldsymbol{\theta},\boldsymbol{Y}) = \pi(\boldsymbol{Y}|\boldsymbol{X},\eta,\boldsymbol{\theta})\pi(\eta)\pi(\boldsymbol{X}|\theta_1,\sigma^2)\pi(\theta_1)\pi(\sigma^2)$ which requires a prior specification for η,θ_1, and σ. In order to reflect an a-priori lack of knowledge we choose vague priors for all parameters. In this model, we use a Half-Cauchy with scale parameter of 25 for σ and a Uniform $(\psi_1^{-1}, \psi_n^{-1})$ where $\psi_(i)$ is the ith largest eigen vector of the spatial neighborhood. As we used a shared-boreder, or rook, neighbor structure wrapped on a torus, the parameter space is (-.25,.25) as each spatial location has four neighbors. The choice of the Half-Cauchy is in line with the recommendations for vague priors for variance components of hierarchical models as outlined in Gelman et al. (2006) and rigorously defended in Polson et al. (2012). We let the prior for η be Uniform(0,1).

Using a gradient descent method with step-halving we found the posterior mode of $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ to be $\sigma^2 = .32$, $\eta = .22$, and $\theta_1 = .20$. Using the z based parameterization described in Section 3.1 we next explored the parameterization $\log \pi(\boldsymbol{\theta}|\boldsymbol{Y})$ and found credible intervals of $\pi(\sigma^2|\boldsymbol{Y}) = (.29,.36), \pi(\theta_1|\boldsymbol{Y}) = (.22,.23)$ and $\pi(\eta|\boldsymbol{Y}) = (.18,.21)$. Fixing $\boldsymbol{\theta}$ at the posterior mode, we then found an approximate 95% credible interval for β_0 to be (.07,-1.67). In further refinement for β_0 was required, we could proceed to use (3.4). This was not done here as β_0 was not the subject of our primary inference.

The posterior maximum and credible interval for σ^2 appear to be slightly lower than expected, but the remaining parameter credible intervals covered

the generating parameter.

Next we simulated from the reaction-diffusion self-excitation model letting , $\beta_0 = 0$, $\alpha = .1$, $\beta = .2$, $\sigma^2 = .25$, and $\eta = .4$

In fitting the model, we again use vague priors for all the parameters. Again, we place a Half-Cauchy prior on σ^2 as described above. In order to conform to the parameter space of α and β, we let $\pi(\alpha) \sim$ Unif $(0, 1)$ and $\pi(\beta|\alpha) \sim$ Unif $(-\frac{\alpha}{2}, 1 - \frac{\alpha}{2})$.

Again using the Laplace approximation technique of section 3, we found the posterior mode of $\pi(\boldsymbol{\theta}|\boldsymbol{Y})$ to be at $\alpha = .085$, $\beta = .19$, $\sigma^2 = .21$, $\eta = .35$. 95% credible intervals for the posterior marginals were $\alpha \in (.07, .10)$, $\beta \in (.14, .24)$, $\sigma^2 \in (.18, .24)$, and $\eta \in (.32, .40)$. At the posterior mode of $\boldsymbol{\theta}$, the posterior marginal for β_0 was approximately (-.03,.01).

In our simulations, the approximations described in this manuscript performed reasonably well for inference on the spatio-temporal diffusion parameters in most cases. However, when σ^2 is large, or when η is large, we have found that the approximations create bias in one or more of the parameters likely due to the high effective number of parameters. However, all approximate likelihood based methods will likely struggle in these cases as well. As noted in Rue, Martino and Chopin (2009), the approximation error in Laplace based methods is related to the number of effective latent variables over the total sample size.

5. Terrorism Data. One region where the reasons for the diffusion of terror and crime still remains unclear is in Iraq during 2003 to 2010. While violence undoubtedly spread throughout the country, it remains unclear how, spatio-temporally, the spread occurred. In Lewis et al. (2012), the authors concluded that self-excitation was present in select cities in Iraq during this time period. The presence of the self-excitation finding was echoed in Braithwaite and Johnson (2015) where the authors also noted a correlation between locations that shared microscale infrastructure similarities.

However, these are far from the only models that could be fit to the data. In a classic work on the subject, Midlarsky, Crenshaw and Yoshida (1980) discuss how heterogeneity between locations can cause correlation in violence or individuals who cause violence can actually physically move from one location to another. The former theory is reflected in the Spatial Correlation in the Spatially Correlated Self-Exciting model and the later theory would correspond to the Reaction Diffusion component of the second model. Midlarsky, Crenshaw and Yoshida (1980) also discuss temporary excitation corresponding to the η parameter in both models. As demonstrated above, these models are all examples of (3.3).

In order to examine these competing theories and demonstrate how Laplace based approximations can be used to fit real world data to models of the class of (3.3) we used data from the Global Terrorism Database (GTD) introduced in LaFree and Dugan (2007). The GTD defines terrorist events as events that are intentional, entail violence, and are perpetrated by subnational actors. Additionally, the event must be aimed at obtaining a political, social, religions, or economic goal and must be conducted in order to coerce or intimidate a larger actor outside of the victim. The majority of lethal events in Iraq from 2005-2010 fit the above category as they were conducted by presumed non-state actors and directed against the United States or the government of Iraq.

The GTD uses a variety of open media sources to capture both spatial and temporal data on terroristic events. The database contains information on what the event was, where it took place, when it took place, and what terrorist group was responsible for the event. From 2003 to 2010 in Iraq, the database contains 6263 terrorist events, the spatial structure is shown in figure 2.

As seen in this map, the majority of the violence is in heavily populated areas such as Baghdad and in the regions north up the Tigris river to Mosul and west through the Euphrates river. In order to model this data, we aggregated the point data to 155 political districts intersected by ethnicities and aggregated the data monthly for 96 months meaning that $\Sigma(\theta)$ in (2.2) is a 14880 x 14880 matrix. Population for each political district was taken from the Empirical Studies of Conflict Project website, available from https://esoc.princeton.edu/files/ethnicity-study-ethnic-composition-district-level.

As the latent process here is likely not zero mean, an intercept was added to the model. The prior for the intercept, β_0 was chosen to be diffuse, such that $\pi(\beta_0) \sim$ Gaus $(0, 1000)$. The log of the population was also added as an offset. The remaining parameter priors were chosen to be vague as well. Specifically, we use a Half-Cauchy with scale parameter of 25 for σ^2 and a Uniform $(\psi_1^{-1}, \psi_n^{-1})$ where ψ are the eigenvalues of $\boldsymbol{H}$. For this neighborhood structure, this corresponded to a Uniform (-.3,.13). Letting $Y(\boldsymbol{s_i}, t)$ be the number of observed events occurring in region $\boldsymbol{s_i}$ at time t, the model is

FIG 2. *Spatial depiction of 6263 events in Iraq.*

$$
\begin{aligned}
(5.1) \qquad & Y(\boldsymbol{s_i}, t)|\mu(\boldsymbol{s_i}, t) \sim \text{Pois } (\mu(\boldsymbol{s_i}, t)) \\
& \mu(\boldsymbol{s_i}, t) = \exp(\beta_0 + \log \text{Pop}(\boldsymbol{s_i}, \boldsymbol{t}) + X(\boldsymbol{s_i}, t)) + \eta Y(\boldsymbol{s_i}, t-1) \\
& X(\boldsymbol{s_i}, t) = \theta_1 \sum_{\boldsymbol{s_j} \in N(\boldsymbol{s_i})} X(\boldsymbol{s_j}, t) + \epsilon(\boldsymbol{s_i}, t) \\
& \epsilon(\boldsymbol{s_i}, t) \sim Gau(0, \sigma^2)
\end{aligned}
$$

The data were fit using both the self-exciting model, (5.1), as well as a non-self exciting model fixing $\eta = 0$ and a non-spatial diffusion version of the self-exciting model fixing $\theta_1 = 0$. All three models were fit using the Laplace approximation outlined in Section 3.1. Note that the addition of covariates modifies the log likelihood for the data model in (3.6). This, in turn, modifies the partial derivatives in (3.8) and (3.9).

Using the Laplace approximation technique, after approximately 30 minutes, the posterior mode was found to be $\left(\beta_0, \theta_1, \sigma^2, \eta, \right)^T = (-14.23, .126, 1.36, .43)^T$. Following the discovery of the mode, the posterior was explored using the z based parameterization described in Section 3.1. Approximate 95% credible intervals were found to be $\beta_0 \in (-14.2, -14.3)$, $\theta_1 \in (.12, .13)$, $\sigma^2 \in (1.2, 1.5)$, and $\eta \in (.38, .46)$. Notably, in fitting this model to the data the effective number of parameters was found to be over 1800, this is quite high and casts doubts on the both the validity of the approximations as well as the ability of the model to accurately describe the data.

We next fit the data to the RDSEM model. The prior for σ remained a half-Cauchy as described above. For the spatio-temporal diffusion parameters we used priors of $\pi(\alpha) \sim$ Unif $(0, 1)$, $\pi(\beta|\alpha) \sim$ Unif $(-\frac{\alpha}{2}, 1 - \frac{\alpha}{2})$, and $\pi(\eta) \sim$ Unif $(0, 1)$. The complete model is the same as (5.1) with the modification that $X(\boldsymbol{s_i}, t) = \frac{\beta}{|N(s_i)|} \sum_{\boldsymbol{s_j} \in N(\boldsymbol{s_i})} X(\boldsymbol{s_j}, t-1) + (1-\beta-\alpha) X(\boldsymbol{s_i}, t-1) + \epsilon(\boldsymbol{s}, t)$.

The Laplace approximation using gradient descent with step-halving took approximately 30 minutes to find the posterior mode at $\left(\beta_0, \alpha, \beta, \sigma^2, \eta\right)^T = (-16.4, .004, .03, .23, .001)^T$. The spatio-temporal diffusion parameters had 95% credible intervals of $\alpha \in (.003, .004)$, $\beta \in (.027, .031)$, $\sigma^2 \in (.19, .27)$, and $\eta \in (0.0008, .004)$. The effective number of parameters for this model was much lower (≈ 800) using this model, suggesting that the approximations valid here.

In comparing the models we calculated DIC as described in Section 3.2, we further calculated posterior predictive P-values based off of the maximum observed value and the number of zeros in the dataset. In the original dataset, the maximum number of events observed for all regions and months was 26 and the dataset had 13445 month/district observations that were zero. The model assessment and selection results are shown in Table 1. From the results we can see that DIC preferred the RDSEM model. The posterior predictive checks suggest that the RDSEM does an adequate job replicating the number of zeros observed in the dataset. However, all models perform poorly in replicating the maximum value observed suggesting additional structure in the models not represented through the baseline structure of the RDSE and the SCSE.

Most significantly, the self-excitation parameter is quite high in the SCSE model whereas it is pushed up against the low edge of the parameter space in the RDSE model. If the SCSEM is used a practitioner may be tempted to conclude that the majority of violence is due to geographic similarities and self-excitation. Whereas, the RDSEM might suggest a policy of isolation to address the movement of violence from neighboring regions overtime.

TABLE 1
Model Assessment and Selection Statistics For Iraq Data

Model	DIC	P-Values Maximum Value	P-Value Zeros
Self-Exciting Model without Spatial Component	10833	1	0
Spatially Correlated Model without Self-Excitation	10865	1	0
Spatially Correlated Self-Exciting Model	10244	.97	0
Reaction Diffusion Self-Exciting Model	8666	.60	.98

6. Discussion. In this manuscript we develop models that allow for spatio-temporal diffusion in the process model and temporal diffusion in the data model. We relate the models to existing theory in how violence diffuses in space and time. We further developed a Laplace approximation for spatio-temporal models that contain self-excitement. This modification allows for a quick and accurate fit to commonly used models in both the analysis of terrorism and criminology. A critical difference between classical INLA and our proposal is that in our proposal, inference is not only performed on the hyperparameters during the exploration of $\pi(\theta|Y)$ in (3.3) but also on the self-excitation parameter η. While η is not a hyperparameter, when the linear expansion of the log-likelihood is done in (3.7), η enters into $\boldsymbol{Q}^*$ and $\boldsymbol{B}^*$ in a similar manner as the hyperparameters.

While we only considered two process models and self-excitation that only exists for one time period, the methodology outlined above can easily be extended to allow self-excitation to have an exponential decay similar to the modeling technique of Mohler et al. (2012). As shown above, the absence of testing multiple process models may result in premature conclusions about how violence is spreading over regions. While self-excitation may be present in one model, its significance may be dulled through the use of alternative process models resulting in differing conclusions.

Although self-excitation has become increasingly popular, alternate approaches based on Besag's auto-logistic model, as used in Weidmann and Ward (2010) are possible, though care must be taken if count data is used as Besag's auto-Poisson does not permit positive dependency. As shown in Kaiser and Cressie (1997), a Winsorized Poisson must be used if positive dependency is desired, as it most certainly is in terrorism modeling. In this case, the data model dependency would linearly be associated with the log of $\mu(\boldsymbol{s_i}, t)$.

As an extension to the analysis of Iraq performed above, it would also be beneficial to add covariates for ethnicity. The choice of ethnicity as a covariate reflects the socio-political theory that spread of terrorism is not only a function of proximity, but similar ethnic groups often emulate each other as in Kuran (1998). Other potential covariates such as road networks that were significant in Braithwaite and Johnson (2015) weren't used as they would be highly collinear with population when aggregated over district. Other macro-scale covariates, such as level of democracy as used in Python et al. (2016) were also not considered as they would be approximately equal in each region.

Though the motivation for the models in this manuscript was the spatio-temporal spread of violence, the novel concept of combining latent process dependency and data model dependency has the potential to be used in other fields. For example in the modeling of thunderstorms, self-excitation may be present temporally, while process model dependency may also be appropriate due to small-scale, unobservable, spatial or spatio-temporal dependency. Laplace approximations, as demonstrated in this manuscript, allow for quick and relatively accurate methods to fit multiple types of self-exciting spatio-temporal models for initial inference.

References.

Besag, J. (1974). Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society. Series B (Methodological)* 192–236.

Braithwaite, A. and Johnson, S. D. (2015). The battle for Baghdad: Testing hypotheses about insurgency from risk heterogeneity, repeat victimization, and denial policing approaches. *Terrorism and Political Violence* **27** 112–132.

Chung, F. R. (1997). *Spectral graph theory* **92**. American Mathematical Soc.

Cressie, N. and Wikle, C. K. (2012). *Statistics for spatio-temporal data*. John Wiley & Sons.

Gelman, A. et al. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian analysis* **1** 515–534.

Gelman, A., Meng, X.-L. and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica sinica* 733–760.

Gelman, A., Carlin, J. B., Stern, H. S. and Rubin, D. B. (2014). *Bayesian data analysis* **2**. Chapman & Hall/CRC Boca Raton, FL, USA.

Kaiser, M. S. and Cressie, N. (1997). Modeling Poisson variables with positive spatial dependence. *Statistics & Probability Letters* **35** 423–432.

Kuran, T. (1998). Ethnic norms and their transformation through reputational cascades. *The Journal of Legal Studies* **27** 623–659.

LaFree, G. and Dugan, L. (2007). Introducing the global terrorism database. *Terrorism and Political Violence* **19** 181–204.

Laub, P. J., Taimre, T. and Pollett, P. K. (2015). Hawkes Processes. *arXiv preprint arXiv:1507.02822*.

Lee, Y. and Nelder, J. A. (1996). Hierarchical generalized linear models. *Journal of the Royal Statistical Society. Series B (Methodological)* 619–678.

Lewis, E., Mohler, G., Brantingham, P. J. and Bertozzi, A. L. (2012). Self-exciting point process models of civilian deaths in Iraq. *Security Journal* **25** 244–264.

Midlarsky, M. I., Crenshaw, M. and Yoshida, F. (1980). Why violence spreads. *International Studies Quarterly* **24** 262–298.

Mohler, G. et al. (2013). Modeling and estimation of multi-source clustering in crime and security data. *The Annals of Applied Statistics* **7** 1525–1539.

Mohler, G. O., Short, M. B., Brantingham, P. J., Schoenberg, F. P. and Tita, G. E. (2012). Self-exciting point process modeling of crime. *Journal of the American Statistical Association*.

Polson, N. G., Scott, J. G. et al. (2012). On the half-Cauchy prior for a global scale parameter. *Bayesian Analysis* **7** 887–902.

Porter, M. D., White, G. et al. (2012). Self-exciting hurdle models for terrorist activity. *The Annals of Applied Statistics* **6** 106–124.

Python, A., Illian, J., Jones-Todd, C. and Blangiardo, M. (2016). A Bayesian Approach to Modelling Fine-Scale Spatial Dynamics of Non-State Terrorism: World Study, 2002-2013. *arXiv preprint arXiv:1610.01215*.

Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the royal statistical society: Series b (statistical methodology)* **71** 319–392.

Short, M. B., D'Orsogna, M. R., Pasour, V. B., Tita, G. E., Brantingham, P. J., Bertozzi, A. L. and Chayes, L. B. (2008). A statistical model of criminal behavior. *Mathematical Models and Methods in Applied Sciences* **18** 1249–1267.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64** 583–639.

Tench, S., Fry, H. and Gill, P. (2016). Spatio-temporal patterns of IED usage by the Provisional Irish Republican Army. *European Journal of Applied Mathematics* **27** 377–402.

Weidmann, N. B. and Ward, M. D. (2010). Predicting conflict in space and time. *Journal of Conflict Resolution* **54** 883–901.

Iowa State University
Department of Statistics
E-mail: nclark1@iastate.edu

pdixon@iastate.edu